

EL 961414558

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**APPLICATION MODEL THAT INTEGRATES THE
WEB EXPERIENCE WITH THE TRADITIONAL
CLIENT APPLICATION EXPERIENCE**

Inventor(s):
Margaret Goodwin
Mark Alcazar

ATTORNEY'S DOCKET NO. MS1-1791US

TECHNICAL FIELD

This invention relates to the creation of software applications, and more particularly, to the development of an application model that integrates the Web experience with the client application experience.

BACKGROUND OF THE INVENTION

Today there are essentially two types of applications: Web applications and traditional desktop or client applications. Web applications are a new type of software program that is commonly deployed as multiple Web pages accessible over the Internet. A conventional Web application includes multiple Web pages representing markup-based documents. The Web application may also include script or other resources that are accessed through the Web pages. Commonly, the Web app is stored on a Web server and downloaded to a local computer when used. Currently, it is not actually possible to download and install a typical Web app for subsequent use when the computer is offline.

Traditional desktop or client applications are software programs that are embodied largely if not exclusively in executable code, rather than HTML pages. Typically, desktop applications are installed directly to a local computer and execute on that computer. They execute with high permissions and typically take full advantage of the platform on which they are installed.

Both types of application have their own strengths and weaknesses. For instance, Web applications are easily installed and typically do not impact the local computer. Web applications provide users with an intuitive mechanism for

1 navigating backward and forwards through the pages of the application. However,
2 Web applications usually cannot take full advantage of the installed platform
3 because they do not have sufficient execution permissions. Users commonly do
4 not like the idea that an application so easily executed has full access to their
5 computers.

6
7 In contrast, users commonly expect the traditional application to have
8 heightened permissions because installing them is such a deliberate process.
9 However, there is no way to limit the permissions under which a traditional
10 application runs, so a user who chooses to install a traditional desktop application
11 has to trust it completely. Also, once a user installs a traditional application,
12 upgrades to the application are more difficult to obtain. Each time the user
13 launches the application, the same version is executed, regardless of whether
14 patches or upgrades to the application are available. Users often accept this
15 tradeoff because the likelihood is significant that upgrading one traditional
16 application may impact or disable another installed application.

17
18 These and other problems have faced software developers for some time.
19 Until now, an adequate solution has eluded those skilled in the art.

20 21 **SUMMARY OF THE INVENTION**

22 The invention is directed at an "Application Model" that brings together the
23 best characteristics of the Web experience with the best characteristics of
24 traditional desktop applications, and bridges the gap between the two worlds. It
25 provides a powerful framework for secure applications, simplifies development,

1 deployment, and maintenance for client applications, and provides a simpler, more
2 consistent user experience.

3
4 The Application Model includes elements that define the scope of an
5 application, its startup and shutdown behavior, and how it manages windows and
6 resources; provide basic navigation functionality, journaling and journal
7 extensibility, browser integration, and Structured Navigation; and define the way
8 an application is deployed, installed, activated, updated, rolled back, and removed
9 from the system in a secure, non-impactful way.

10 11 **BRIEF DESCRIPTION OF THE DRAWINGS**

12 Figure 1 is a functional block diagram generally illustrating an "Application
13 Model" that brings together the best characteristics of the Web experience with the
14 best characteristics of traditional desktop applications, and bridges the gap
15 between the two worlds.

16
17 Figure 2 is a functional block diagram of a network environment having a
18 client computer, a server, and an application constructed in accordance with the
19 Application Model.

20
21 Figure 3 is a functional block diagram generally illustrating certain
22 components of an application implementing core elements of the Application
23 Model.

1 Figure 4 is a functional block diagram illustrating components of a system
2 that implements security in accordance with the Application Model.

3
4 Figure 5 illustrates a computing device that may be used in illustrative
5 implementations of the present invention.

6 7 **DETAILED DESCRIPTION**

8 Figure 1 is a functional block diagram generally illustrating an "Application
9 Model" 100 that brings together the best characteristics of the Web experience
10 with the best characteristics of traditional desktop applications, and bridges the
11 gap between the two worlds. It provides a powerful framework for secure
12 applications, simplifies development, deployment, and maintenance for client
13 applications, and provides a simpler, more consistent user experience. The
14 Application Model includes three elements: an Application Framework 101, a
15 Navigation Framework 103, and Application Lifecycle Management 105. Each of
16 these elements is described briefly here, and described in detail below.

17
18 The Application Framework 101 defines the scope of the application, its
19 startup and shutdown behavior, and how it manages windows and resources,
20 whether local or online. The application framework 101 provides the power and
21 extensibility of the traditional desktop application in a way that is accessible to
22 Web developers. It also supports a new genre of dialog that is more effective and
23 less annoying than traditional modal dialogs and better integrated with its
24 environment.

1 The Navigation Framework 103 builds on the Application Framework 101
2 to bring the best of the Web experience to the client application space. It includes
3 basic navigation functionality, journaling and journal extensibility, browser
4 integration, and Structured Navigation, which is a new and better way to build
5 wizards and applications to walk users through complex or unfamiliar procedures
6 in a way that leverages the familiar look and feel of the Web.

7
8 Application Lifecycle Management 105 defines the way an application is
9 deployed, installed, activated, updated, rolled back, and removed from the system
10 in a secure, non-impactful way that is easy to administer and less intrusive to
11 users, similar to Web deployment today.

12 13 **ILLUSTRATIVE IMPLEMENTATIONS**

14 Figure 2 is a functional block diagram of a network environment having a
15 client computer 201, a server 202, and an application 210 constructed in
16 accordance with the Application Model 100 being described in this document.
17 Currently, the application 210 is stored on the server 202 and is available for
18 download to the client computer 201.

19
20 The application 210 includes an application manifest 211, an application
21 definition file 212, and a deployment manifest 213. The application may also
22 include other files (not shown) containing markup, code, and resources belonging
23 to the application 210. The application manifest 211 declares the resources
24 belonging to the application, as well as security permissions being requested by
25 the application 210. The security permissions are discussed below in conjunction

1 with Figure 4. The application manifest 211 also describes the application's
2 dependencies, such as additional libraries and resources that are needed to run the
3 application 210. The application manifest 211 also specifies shell integration
4 information, like file associations for the application 210 and whether to create
5 shortcuts on the client computer 201 in the Start menu.

6
7 The application definition file 212 provides a declarative definition of the
8 application 210. The application definition file 212 may be written in XAML
9 (eXtensible Application Markup Language) and code (using any managed code
10 language). The deployment manifest 213 includes a URI that points to a current
11 version of the application 210, and the intervals at which to check for updates.

12 13 **Application Lifecycle Management**

14 Application Lifecycle Management refers to the characteristics and features
15 of the Application Model that relate to managing the installed state of an
16 application throughout its lifecycle. To begin, the application 210 can either be
17 run from the sever (as an "Express Application"), or locally (as a "Managed
18 Application"). In most cases, these choices don't require any code changes. They
19 simply require changing an attribute in the application manifest 211.

20
21 An Express application is an application that runs from the server 202.
22 While an Express application may be cached on the client computer 201 in a
23 transient cache 250, it may be scavenged from the cache 250 at any time, based on
24 available resources and on how recently it was used, so it is not guaranteed to
25 work offline. An Express application may also require online resources to run

properly. Express applications are always hosted in the browser, and cannot request permissions beyond a default permission set. This ensures that the user can navigate to the application and it loads without any prompts or dialogs, just like a Web page. Express applications don't create shortcuts or aliases to themselves, and have no permanent impact on the client computer 201. The advantage of Express applications over Web pages is that they are real applications that can take full advantage of the Application Framework and the Navigation Framework (described below), can run locally so they have better performance than server-side applications, and can fully exploit the client computer hardware and operating system features.

Managed applications are deployed the same way as Express applications. The application provider simply copies the application directory 215 and the deployment manifest 213 to the server. The user also accesses Managed applications the same way they access Express apps. The user navigates to an URL that points to the deployment manifest 213. Like Express applications, Managed applications download progressively to give the user immediate interaction with the application. Unlike Express applications, Managed applications are available offline, and are integrated with the client computer environment, which makes them more discoverable. Examples of this integration included shortcuts or aliases to, and file associations for, the Managed applications. Once a managed application is installed, it can either be launched from the Shell using a shortcut or alias, or by double-clicking a file associated with the application, or it can be navigated to in the browser by clicking a

1 hyperlink, entering a URI in the address well, or selecting a stored hyperlink in a
2 Favorites menu. Managed applications are always launched and run locally.

3
4 An installer 252 uses the application manifest 211 to control download and
5 installation. The installer 252 also invokes a trust tanager 254 during installation.
6 The trust manager 254 uses the application manifest 211 to determine what
7 permissions the application requires to run. (See Security, below), what file
8 associations to register, and the application name and icon to use in the Start menu
9 and task bar.

10
11 When the user first navigates to the application 210, it is actually the
12 deployment manifest 213 they navigate to. The deployment manifest 213 points
13 to the URI where the latest version of the application is published, and the
14 installer 252 downloads and installs the application 210 from there. This makes it
15 possible to post updates to the application 210 without disrupting any downloads
16 currently in progress, while keeping the same URI for the application 210. When a
17 new version is posted, the deployment manifest 210 is changed to point to the new
18 version. Any downloads in progress continue and any new navigations to the
19 deployment manifest 213 are redirected to the latest version.

20 21 ***Progressive installation***

22 One of the advantages of the Web over traditional desktop applications is
23 that when a user navigates to a Web page or application, it downloads
24 progressively, so they can start interacting with it immediately. Traditionally, that
25 hasn't been possible with desktop applications. The Application Model supports

1 progressive download for Managed applications as well as Express applications.
2 The application manifest 211 specifies subsets of components associated with the
3 application 210 as Required, OnDemand, and Online. The subset designated as
4 Required should include the minimum code necessary for the application 210 to
5 run on the client computer 201, and will always include at least the application
6 definition file 212 and the application's startup page. The subset of the application
7 210 designated as Required is downloaded immediately and stored in an
8 application cache 260 on the client computer 201, and the application 210 is
9 launched. Because the minimum code is considerably smaller than the full
10 application, the user can begin interacting with the application right away, similar
11 to a user's experience when interacting with a traditional Web application.
12

13 As the user interacts with the application 210, the resources designated as
14 OnDemand are drizzle-download in the background. When the user requests a
15 specific resource, e.g., by clicking a hyperlink, that resource takes precedence over
16 the other resources that are drizzling down in the background and is downloaded
17 on demand. As additional OnDemand resources are downloaded, they're stored in
18 the application cache 260 as well. Resources that are designated as Online
19 resources are stored in the transient cache 250. As resources are downloaded, the
20 next time they're requested, they're retrieved from the application cache 260.
21 Once the application 210 has been fully installed on the client computer 201, it
22 may be navigated to in the browser or invoked from the shell, but it is always
23 launched locally from the client computer 201.
24

25 *Update, Rollback, Uninstall*

1 When downloading the application 210, the installer 252 downloads a copy
2 of the deployment manifest 213, which specifies the Uri to check for application
3 updates and the intervals at which to check. The installer 213 polls the server 202
4 for updates at the intervals specified in the deployment manifest 213 to determine
5 whether there is an update available. If there are any updates, the updated
6 components of the application 210 are downloaded transparently in the
7 background without interrupting the user. This ensures that the user always has the
8 latest version of the application 210 whenever they launch it.

9
10 Users can painlessly rollback or uninstall the application 210 through
11 Add/Remove Programs, and it will not impact other applications on the client
12 computer 201. This is because applications are installed in their own isolated
13 cache, with all the components they need to run. An application will not overwrite
14 a resource required by another application because it has its own copy of all its
15 dependencies. Because shortcuts and file associations are declared in the
16 application manifest 211, which is used for installation, rollback, and uninstall,
17 these will automatically be removed if the application 210 is uninstalled, and will
18 be mapped to the appropriate version if the application 210 is rolled back.

20 *Host in Browser vs. Standalone*

21 One of the unique aspects of the Application Model is the ability to write
22 the same application and decide at compile time whether to host it in the browser
23 275 or a stand-alone top level window 276. An attribute in the application file,
24 named HostInBrowser, is used for this determination. Setting this attribute to
25 “True” specifies that the application 210 should be hosted in the browser 275.

1 Setting it to "False" specifies that the application 210 will be hosted in its own
2 top-level window 276.

3
4 Unlike existing technologies, the application code is essentially identical in
5 either case. Applications execute in a common execution environment 280 that
6 can be either owned by the browser or not. When the code is launched, a host
7 sniffer component 281 determines from the HostInBrowser attribute which hosting
8 environment is appropriate. Based on that indicator, the code is launched in the
9 appropriate hosting environment.

10
11 Launching the application 210 can occur in two ways: either by activating a
12 link to the locally stored application, or by navigating to an URL in the browser.
13 When launching the application 210 from the link, the shell may launch a shell
14 extension handler registered to handle files of the type that was activated. The
15 Shell extension handler for Application Model applications uses a host sniffer to
16 determine the hosting environment in which to invoke the code. The host sniffer
17 determines the value of the HostInBrowser attribute, and either instructs the shell
18 server to launch the application 210 or to invoke the browser 275 with instructions
19 to navigate to it.

20
21 Likewise, navigating to the application 210 using the browser 275 makes
22 use of a mime-type handler, which is a handler registered with the browser 275 to
23 handle mime-types associated with applications that may be hosted in more than
24 one environment. The browser server proxy is configured to determine the host
25

environment for the application using the host sniffer 281 and launch it in the appropriate hosting environment.

Today, whether an application runs in the browser or standalone is determined by whether it's run from a Web server or the local machine. Ideally, the hosting environment would be determined by the requirements of the application, rather than where the bits are stored.

The browser 275 is well-suited to task-based applications, whose purpose is to accomplish a specific task, like retrieving, submitting, or tracking data, configuring or managing resources, and the like. Hosting in the browser provides convenient serial access to applications, and relieves the user of having to manage multiple windows.

Standalone windows 276 are well suited to background-style applications, whose purpose requires that they stay open while the user interacts with other applications in the browser, and to applications that are not based on the navigation model. For example, document/content editors don't gain anything from being navigation-based. In this type of application, actions are initiated by the user, and there is no common sequence in which actions occur.

Background-style applications may be hosted in a standalone window, and can be minimized when the user isn't directly interacting with them. For example, people use media players to play music while they work, and they frequently interact with them to select a new song or playlist. This type of interaction should

1 not interrupt their workflow. In other words, they shouldn't have to navigate away
2 from the application they're currently using to perform some task in order to select
3 a new song or playlist (especially if navigating to a new playlist would destroy
4 their Forward stack so they couldn't easily return to their task). These types of
5 applications make sense as a navigation application, but do not make sense to host
6 in the browser 275.

7
8 The Application Model is equally well suited to either type of hosting, but
9 the decision about how to host a particular application is up to the developer.
10 Today, a developer has to decide before writing code whether the application will
11 be a Web application or a traditional desktop application. This decision dictates
12 everything from what language and tools the developer can use to write the
13 application to the environment in which the application will be hosted at runtime.
14 Because different types of applications require different tools and skillsets to write
15 them, developers are typically limited to developing either Web applications or
16 desktop applications, but seldom both. This invention makes it possible to use the
17 same skillset and tools to develop either type of application.

18
19
20 Figure 3 is a functional block diagram generally illustrating certain
21 components of an application 301 constructed in accordance with the Application
22 Model 100. The application 301 is essentially illustrated as a group of interrelated
23 classes of code that may be used to implement the application 301. It will be
24 appreciated that, when executed, many of the classes described will be instantiated
25 as objects, although certain properties, methods, and events of the classes may still

1 be employed statically. The components illustrated in Figure 3 service both the
2 Application Framework and the Navigation Framework elements of the
3 Application Model. Accordingly, those two elements will be described with
4 reference to the components illustrated in Figure 3.

6 **Application Framework**

7 In traditional desktop programming, there is a Main function that defines
8 the lifetime of a desktop application. The desktop application begins when Main()
9 is invoked, and ends when Main() returns. In conventional Web programming,
10 there is no corresponding function or object that scopes the application. In this
11 Application Model, the application 301 is defined and scoped by an application
12 object, which is a global object that persists in memory for the lifetime of each
13 application session. The application object is an instance of an application
14 class 303, described below. The application object has knowledge of the resources
15 belonging to the application 301, and provides an isolation boundary between
16 itself and other applications or external resources. It is used by the platform to
17 identify, reference, and communicate with the application 301. It is also used
18 within the application 301 to manage windows and resources, specify startup and
19 shutdown behavior, handle configuration settings, specify visual styles for the
20 application, share code, state, and resources across navigations, and handle
21 application-wide events.

23 ***Application class***

24 The basic application class 303 is intended to be able to support any
25 number of different types of applications, so it is designed to be as low-policy as

1 possible. The only methods on the basic application class 303 are ShutDown,
2 which closes the application 301, FindResource, which can be used to retrieve any
3 resource belonging to the application 301, and GetService, which is used to
4 retrieve any service exposed by the application 301.

5
6 On execution, the application object raises events on StartingUp,
7 ShuttingDown, Activate (raised when one of the application's windows receives
8 focus, after another application has had focus) and Deactivate (when the
9 application loses focus), and SessionEnding (raised when the application is being
10 ended by the system either because the system is shutting down or the user is
11 logging off). The application can handle these events to perform whatever actions
12 are necessary in response to the condition that raised them. The StartingUp,
13 ShuttingDown, and SessionEnding events are cancelable, allowing the developer
14 to prevent the action from taking place. In one embodiment, canceling
15 SessionEnding is treated as a hint. There is no guarantee the session will not end,
16 but it gives the developer a chance to save state or perform other necessary
17 cleanup activities before the application is terminated.

18
19 The application class 303 has a static property of type "Application," called
20 Current, that returns the current instance of the Application class 303. The
21 Application class 303 has a Windows collection, which enables the application
22 301 to manage the windows belonging to the application 301. Whenever a new
23 window is created, it is automatically added to this collection. The Application
24 class 303 also has a MainWindow property which is used to get the application's
25 primary window. For an application hosted in the browser, this may be the

1 browser window. For other applications, this is set by default to the first window
2 created, but may be explicitly set to a different window by the developer.

3
4 The Application class 303 also has a Resources collection, which is used to
5 store resources, such as styles. Styles are rules for visually styling elements in a
6 tree hierarchy. The scope of a style is the element to which the style is applied and
7 its children. Styles are defined as resources so they can be added to the Resources
8 collection. When a style is specified at the application object level, it applies to all
9 the windows within the application 301 and, therefore, to every visual element in
10 the application 301. A developer can use application object-level styles to enforce
11 a consistent look and feel throughout the application 301.

12
13 The Application class 303 also defines a Settings property that is used to
14 store and retrieve application settings. Finally, it has a ShutDownMode property,
15 which controls what action causes the application 301 to be terminated. The
16 options may include: Shut down when the Main Window is closed, Shut down
17 when the last window has been closed, or only shut down when the ShutDown
18 method is called.

19 20 *Window class*

21 Windows in the Application Model support features like transparency,
22 rotation, non-rectangular surface, etc., which aren't possible with conventional
23 desktop windows today. The Window class 305 provides a service, by invoking
24 GetService, that exposes a conventional window handle for the Window that can
25

1 be used for interoperability with legacy windows, when parenting a legacy
2 window to an Application Model window and vice versa.

3
4 The Window class 303 provides the typical properties available on
5 windows today for accessing title, icon, size, position, style, state, whether the
6 window has focus, whether it is visible, whether it is closed, whether to display it
7 in a taskbar, whether it should be topmost, accessing its owned windows, etc. It
8 also has a DialogResult property for obtaining the result if the window is a dialog.

9
10 The Window object raises events when it is Loading (about to open),
11 Closing, Activated (receives focus), Deactivated (loses focus), and when its size,
12 location, or state is changing.

13
14 The Window class 305 has methods to Show, Hide, Close, DragMove,
15 GetService, ShowDialog, and ShowApplicationDialog. ShowApplicationDialog is
16 equivalent to the conventional ShowDialog method, which shows the standard
17 type of dialog users are familiar with today. This is a floating thread-modal dialog
18 that blocks interaction with any window on the thread from which it was opened
19 until the user closes it.

20
21 In contrast, the ShowDialog method of the window class 305 displays a
22 new kind of dialog, which is modal only to its parent window. This dialog blocks
23 interaction with the parent window's content, but other windows in the application
24 301 are still interactive. The dialog's parent window can still be moved, resized,
25 minimized, maximized, or restored while the dialog is open. If the parent window

1 is move or minimized, the dialog is moved or minimized along with it. If the
2 parent window is minimized and then restored, the dialog is restored along with it,
3 and the window content is still disabled until the user closes the dialog. This is a
4 new concept for programmers and will provide a much better user experience for
5 most cases when a dialog only applies to the content of a single window because it
6 allows users to continue interacting with other windows in the application in a
7 normal manner.

9 ***Pane class***

10 The pane class 307 provides a layer of abstraction that makes it possible to
11 programmatically access common window properties without having to know
12 whether the “window” is an actual stand-alone window or the system browser.
13 The pane class 307 gets the BrowserWindowService, if hosted in the browser, or
14 the IWindowService, if hosted in a standalone window, and uses that to
15 communicate with the window or browser without the application 301 having to
16 know which one is actually present. The pane class 307 exposes a Navigator
17 property (the INavigator of the containing Window or BrowserWindowService),
18 as well as all of the IWindowService properties (described below).

20 ***IWindowService interface and the BrowserWindowService class***

21 The IWindowService interface 309 exposes the following properties: Text
22 (the title), Icon, Top, Left, Height, Width, Maximum/MinimumContentHeight,
23 Maximum/MinimumContentWidth, ContentSize, WindowState (Maximized,
24 Minimized, Normal), and WindowAutoLocation (CenterScreen, CenterOwner,
25 WindowsDefaultLocation, None).

1
2 The BrowserWindowService class 311 provides the properties exposed by
3 IWindowService and, in addition, exposes the following properties and methods
4 that are unique to the browser: DisplayMode (Normal, FullScreen, TheaterMode)
5 and ShowMenuBar. It also has methods to start and stop graphical indications that
6 a navigation is in progress, and a GetService method. A developer that wants to
7 access browser-specific functionality can obtain the BrowserWindowService 311
8 by invoking GetService on the application object or on a NavigationWindow or
9 Pane, providing that the application 301 is hosted in the browser. If it is not
10 hosted in the browser, the request will return null.

11
12 Using the BrowserWindowService 311 explicitly sacrifices the host-
13 agnosticism that allows the application 301 to be hosted either in the browser or
14 standalone without making any significant change to code.

15 16 **Navigation Framework**

17 The third and final element of the Application Model 100 is the Navigation
18 Framework. The Navigation Framework supports navigation-based applications
19 that leverage users' familiarity with navigation and journaling metaphors on the
20 Web to provide a more familiar, consistent user experience on the client, whether
21 hosted in the system browser or a stand-alone top level window.

22 23 *NavigationApplication class*

24 The purpose of the NavigationApplication class 315 is to make it easier for
25 developers to write navigation-based applications. It is derived from the

1 Application class 303, so it contains the properties, methods, and events exposed
2 by that class. Additionally, the NavigationApplication class 315 has a StartupUri
3 property that is used to specify a page or element to navigate to when the
4 application 301 is first launched. It also has a Properties collection that allows the
5 application developer to share global state across pages without having to subclass
6 the Application, and supports databinding to these properties. A Connected
7 property tells whether the machine the application is running on is currently
8 connected to a network.

9
10 The NavigationApplication class 315 also listens for navigation events on
11 NavigationContainers (described below) within the application 301 and re-fires
12 the events, so the application developer can create handlers for all application-
13 wide events in a single place. If a NavigationApplication is not hosted in the
14 browser, the default window created for it is a NavigationWindow.

15
16 The NavigationApplication object provides a Properties collection to store
17 and retrieve arbitrary application state that can be shared across the entire
18 application as well as providing a single place to handle navigation events
19 throughout the application 301. By using the NavigationApplication class,
20 properties can be stored without having to learn subclassing. This simplifies the
21 development process for developers who aren't familiar with the concept of
22 subclassing.

23 24 *INavigator interface and NavigationContainer class*

25

1 The INavigator interface 308 exposes navigation-related properties,
2 methods, and events for the application 301. Any element that allows navigation
3 implements the INavigator interface 308. The INavigator interface 308 provides
4 an ID property to allow it to be specified as the target of a hyperlink navigation. It
5 has a Uri property that can provide the current Uri of the INavigator's content if a
6 navigation is not in progress, or the Uri being navigated to if there is a navigation
7 in progress. Setting this property initiates a navigation to the Uri specified. There
8 is also a CurrentUri property, which is read-only, and returns the current Uri, even
9 when a navigation is in progress. A Content property provides access to the root
10 of the element tree contained within the INavigator. A Journal property returns
11 the Journal instance in which navigations are tracked, while CanGoForward and
12 CanGoBack properties tell whether there are any entries in the Forward or Back
13 branch of the Journal, respectively. A Synchronous property specifies whether
14 navigations are performed synchronously or asynchronously. (See
15 *Synchronous/Aysynchronous Navigation* below.)

16
17 The INavigator interface 317 has several overloaded navigation methods,
18 which variously allow navigation to a Uri or root element of an element tree, and
19 allow the developer to specify whether the navigation is synchronous or
20 asynchronous. It is also possible to pass arbitrary data (as an object) along with the
21 navigation. This data can be retrieved in the navigation event handlers. The
22 INavigator interface also exposes a Refresh method to refresh the content by
23 reloading it from the source, a StopLoading method to cancel any pending
24 download operation, and a GoForward and GoBack method to initiate a Journal
25 navigation.

1
2 INavigator interface 317 implementations raise events for stages of
3 navigation to allow complete flexibility for the developer to modify the behavior
4 to suit the application 301. Navigating (fired before a navigation begins),
5 LoadStarted (fired immediately after a navigation begins), NavigationProgress
6 (fired periodically throughout the navigation to indicate how much has been
7 downloaded and how much remains), NavigationError (fired when an error occurs
8 in navigation or download), Navigated (fired when the target has been found and
9 download has begun), and LoadCompleted (fired when all content has been
10 downloaded and parsed) are examples. Of these, only the Navigating event can be
11 canceled, which prevents the navigation from occurring.

12
13 The NavigationContainer class 319 provides an implementation of
14 INavigator that is used by Frame. It can also be reused by third party component
15 developers who want to implement new types of navigable elements or classes. In
16 addition to INavigator, it also exposes IJournalData, which has methods to enable
17 persisting state to the journal when navigating away from the current page, and
18 restoring state when navigating back to it by way of the Journal (Back/Forward
19 navigation).

20 21 *NavigationWindow class*

22 The NavigationWindow class 321 is derived from the Window class 305,
23 but also implements INavigator by delegating to a Frame element in its Visual
24 Tree. Both Window and NavigationWindow support styling of their Visual tree.
25 With NavigationWindow, styling enables part of the window to persist across

1 navigations so that “navigation chrome” (back / forward buttons) remain as
2 navigations occur.

4 **Synchronous/Asynchronous Navigation**

5 On the Web today, navigation is asynchronous. Some applications,
6 particularly when installed locally, may prefer synchronous navigation so that UI
7 on a page appears simultaneously. There are also times when a developer may
8 want a particular navigation instance to be synchronous, for example, when
9 navigating to a PageFunction (described below). If the developer wants to pass
10 some state to the new PageFunction that will be reflected in the UI the
11 PageFunction displays, he needs to be able to navigate to the UI and then set a
12 property on it in the PageFunction’s constructor. If the navigation is asynchronous,
13 he can’t do this because the element he wants to change may not have been parsed
14 yet. The developer could use the LoadCompleted event for this, but the relevant
15 element may already be displayed before it gets changed. Accordingly, the
16 Application Model provides mechanisms to enable synchronous navigation.

17
18 By default, navigation is asynchronous unless otherwise specified.
19 However, a developer may specify that an application’s default navigation
20 behavior is synchronous by setting the value of the Synchronous attribute in the
21 application class 303 definition to “true”. The developer can also specify that the
22 default navigation behavior for a particular INavigator (NavigationContainer or
23 NavigationWindow) is synchronous by setting the Synchronous attribute on the
24 NavigationContainer object or NavigationWindow object to “true”. If the default
25 navigation behavior for the application is synchronous, a developer can override

1 that for a specific NavigationContainer object or NavigationWindow object by
2 setting the Synchronous attribute on it to "false." To specify synchronous
3 navigation on a hyperlink, when the application or containing INavigator is not
4 synchronous, a developer can set the value of the Synchronous attribute on the
5 hyperlink to "true" (or vice versa).

6 7 **Journaling**

8 Journaling is the process by which the Navigation Framework tracks
9 history. A "journal" enables users to retrace their steps backward and forward in a
10 linear navigation sequence. Whether a navigation application is hosted in the
11 browser or in a standalone NavigationWindow, each navigation is persisted in the
12 journal, and can be revisited in a linear sequence by using the forward and back
13 buttons or by invoking the GoForward and GoBack methods. Each
14 NavigationWindow has its own Journal object.

15 16 ***Journal class and JournalEntry class***

17 The default behavior of the Journal is the same as the Journal behavior of
18 the browser today, thereby providing users with a familiar Web experience. In
19 addition, the application 301 can insert virtual Journal entries that perform DOM
20 operations on the current page, and applications can remove entries from the
21 journal, but only within the application's own journal (not the system journal).

22
23 The Journal class 325 has the following methods: AddEntry, RemoveEntry,
24 and GetEnumerator. GetEnumerator returns an enumerator over the current list of
25 JournalEntries in the Journal. The RemoveEntry method takes an offset and

1 removes the JournalEntry at that that offset in the Back stack. The overloaded
2 AddEntry method allows the developer to insert a JournalEntry either at the
3 current position, or at a specified offset. To add a journal entry to the journal, an
4 application derives a class from the JournalEntry class 327.

5
6 The JournalEntry class 327 has two properties, Name and Uri. The name is
7 the text that is displayed in a dropdown list on the Back/Forward buttons, and the
8 Uri is the URI of the stored page. The JouranlEntry class 327 has only one
9 method, called Replay. When adding a custom entry to the Journal, the derived
10 class overrides the Replay method to perform arbitrary DOM operations on the
11 tree to restore it to its previous state. The derived class should contain whatever
12 properties or methods are necessary to store and replay the state changes. The
13 overridden Replay method returns a new JournalEntry that undoes whatever action
14 it performed in the method itself. This JournalEntry will be inserted in the Journal
15 to be replayed when the inverse Journal navigation is performed.

16
17 For example, if the custom JournalEntry is being replayed as the result of a
18 Back navigation, the JournalEntry returned by its Replay method is placed in the
19 Forward stack. When the user navigates Forward, this journal entry is replayed,
20 restoring the state of the page just prior to the Back navigation. Likewise, when a
21 Forward navigation takes place, the returned JournalEntry is placed in the Back
22 stack, so when the user navigates Back, the state of the page can be restored. The
23 JournalEntry returned by the Replay method should perform the inverse of
24 whatever operation the Replay method performed.

1 Classes derived from `JournalEntry` should implement `ISerializable`, so they
2 can be stored in the system `Journal` if the application is hosted in the browser.

4 ***IJournalData* class**

5 When a navigation takes place, the controls on the page that is going away
6 are given an opportunity to persist their state to the journal, and to reload it when
7 the user revisits the page. To persist its state in the journal, and restore it when the
8 page is revisited, a control implements the `IJournalData` interface 323. The
9 `IJournalData` class has two methods, `Save` and `Load`, and a `PersistID` property that
10 is used by the `Journal`.

12 **Structured Navigation**

13 In many traditional applications, where the user is either familiar with the
14 task, or where the user is creating arbitrary content, it is natural for the user to lead
15 the application. For infrequent or complex tasks that require specific procedures to
16 be followed in a specific order, it makes more sense for the application to lead the
17 user. This is best accomplished by providing a structured navigation sequence
18 through pages that each have a focused purpose, communicated to the user
19 through a direct question or instruction, and a simple control layout so the user
20 knows exactly what's expected at any given point. Structured navigation is well
21 suited to this kind of application. The `PageFunction` class 329 is the building block
22 for Structured Navigation.

24 Wizards tend to provide a single path of navigation through a series of
25 steps, while Web pages generally allow arbitrary navigation. However, even

1 many Web sites require a specific path for certain tasks. For example, a checkout
2 line in a shopping application. Also, many wizards would be more effective and
3 less frustrating to users if they could provide a non-linear navigation sequence,
4 based on current state. For example, a wizard that requires an Internet connection
5 to complete its main purpose would be more effective if it could detect when a
6 user is already connected to the Internet, and bypass the steps to connect to the
7 internet without having to present the user with that decision.

8
9 A PageFunction is like a Web page because it is invoked by a navigation
10 request, and it can display UI like any other page. It is like a function because
11 each PageFunction object encapsulates a unit of functionality. Its caller can pass
12 state to it, and it can return a result when it has completed its purpose.
13 PageFunctions can be nested and chained. A PageFunction can also invoke other
14 PageFunctions and delegate to them to carry out all or part of its functionality. A
15 PageFunction can determine the next action to take (or the next PageFunction to
16 navigate to) based on accumulated state. A PageFunction can simply be a decision
17 point that delegates to other PageFunctions, but has no UI itself. This allows for a
18 very flexible navigation model, in which the user can be led through a complex
19 procedure, omitting any steps that aren't relevant.

20
21 A PageFunction can also specify that once it completes, it and all of its
22 children (other PageFunctions it invoked in the process of completing its purpose)
23 should be removed from the Journal. An example of where this might be useful is
24 when a user has completed a transaction, or made a purchase. For security
25 reasons, it would be best if their credit card information isn't still available if

1 someone should click the Back button. With the PageFunction class 329, the
2 entire transaction can be removed form the Journal after the user reaches the
3 confirmation page, so when the user clicks Back, he's returned to the point in the
4 application before the transaction was started. Clicking Forward would return to
5 the confirmation page.

6 7 *PageFunction class*

8 The PageFunction class 329 uses Generics (similar to Templates in C++).
9 This provides a way for the same class to be reused with different parameters in a
10 strongly typed way. In the case of PageFunctions, the parameter represents the
11 ReturnArgsType 330, which is the type of the value returned by the PageFunction
12 on completion.

13
14 The PageFunction class 329 has two properties, RemoveFromJournal,
15 which specifies whether the PageFunction and its children should be removed
16 from the Journal after completion, and KeepAlive, which specifies whether the
17 PageFunction should be kept in memory after navigating to a child PageFunction.
18 By default, the parent is "dehydrated" when navigating to a child, and
19 "rehydrated" when the child returns. This helps keep the working set down.
20 However, some PageFunction classes are expensive to initialize so, in these cases,
21 it may be worth the extra working set to not have to recreate the PageFunction
22 whenever a child returns. This is an implementation decision that will be available
23 to developers.

1 The PageFunction class 329 has two methods. A virtual Start method is
2 overridden to initialize the PageFunction the first time it is navigated to. This
3 method is not called when a child PageFunction returns, only when the
4 PageFunction is invoked. A virtual OnFinish method is called by the
5 PageFunction itself when it is ready to return to its caller. The parameter is also a
6 Generic class that takes the same type as the Generic parameter for the
7 PageFunction class. This is the value that is returned to the PageFunction's caller.

8
9 The PageFunction class 329 also has an event, called Return. Whoever
10 invokes the PageFunction attaches a listener to this event. When the PageFunction
11 returns, the caller gets the return value from the EventArgs for this event, which is
12 parameterized with the same type as the PageFunction itself.

13 14 **Localization**

15 All resource access is done through relative URIs. When a resource is
16 loaded, the URI is resolved to an absolute URI using the codebase of the
17 application. Because URIs are resolved at run-time, resource requests can be
18 remapped at run-time based on localization settings. For example, the codebase
19 for the application `http://myapp`, a request for `pages/page1.xaml` may be found
20 in the following locations, depending on the current localization setting:

- 21 • `http://myapp/pages/enu/us/page1.xaml`
- 22 • `http://myapp/pages/enu/page1.xaml`
- 23 • `http://myapp/pages/page1.xaml`

1 The Application Model allows for UI to be localized by replacing the entire
2 markup associated with the UI. Since resource loading is done through the same
3 mechanism, it is possible to localize any type of resource, not just string resources.
4

5 **Security**

6 Figure 4 is a functional block diagram illustrating components of a system
7 that implements security in accordance with the Application Model. Today,
8 conventional desktop applications run with full trust. Once an application is
9 installed, it can do anything that the user can do. Web applications, on the other
10 hand, run with a restricted permission set. This helps protect the user from
11 malicious or buggy applications that could do harm to the users machine or data.
12 However, some legitimately safe applications require permissions beyond the
13 restricted permission set to be able to carry out their intended functionality. In
14 accordance with the Application Model, applications don't get full trust just
15 because they're installed locally.
16

17 An application 410 in accordance with the Application Model runs in a
18 Secure Execution Environment (SEE) 420. The SEE 420 defines a default
19 permission set 425, which is more limited than that experienced by the typical
20 desktop application installation. The SEE 420 may interact with a hosting
21 environment 440 that provides policy-based enforcement of access rights to local
22 resources. When the application 410 attempts to access a resource 414, the
23 application's identity and site of origin are verified, and a check is performed to
24 determine whether the application's permission grant set 411 allows it to access the
25

1 requested resource 415. If so, then the requested access is allowed. If not, then
2 the request is denied.

3
4 By default, applications run with the default permission set 425.
5 Applications can request further permissions, if required, by specifying the
6 permissions they need in the application manifest 412. If an application requires
7 permissions beyond the default permission set, installation is required. At
8 installation time, a Trust Manager 421 helps users decide if the application is
9 trustworthy. At installation time, the Trust Manager 421 inspects the application
10 manifest 412 to determine whether the application 410 is requesting any
11 permissions beyond the default permission set 425. If so, the Trust Manager 421
12 evaluates the risk level of the permissions being requested, along with other
13 factors, like whether the application 410 is signed, whether it meets logo
14 requirements, etc., and determines an overall risk rating for the application 410.
15 The Trust UI that is displayed to the user is different depending on the risk rating,
16 to help the user make an informed decision. If the user doesn't grant the
17 permissions requested, the application may not be installed or allowed to run. If
18 the user accepts the heightened level of permissions requested, the application's
19 permissions are stored in association with the application as the application's
20 permission grant set 411.

21
22 The default permission set 425 is intended to allow maximum functionality
23 without compromising the security of the system. For example, an application that
24 needs to store configuration settings can use certain configuration APIs with the
25 default permission set 425 that guarantee no application can overwrite or modify

1 another application's settings. The default permission state may allow an
2 application to store state other than settings, for example, any content or state that
3 should not be rolled back when the application is rolled back. The default
4 permission set 425 may allow the application 410 to launch File Open/Save
5 dialogs. The File Open/Save dialogs provide a means for the user to store/retrieve
6 content in a user-selected file location, without making this information available
7 to the application 425. These activities are considered safe because they don't
8 allow the application 410 direct access to the file system. However, if an
9 application requires direct programmatic access to the file system, it must request
10 File I/O permission, in which case, the user should be prompted by the Trust
11 Manager 421 when the application 410 is installed.

13 **Illustrative Computing Environment of the Invention**

14 Figure 5 illustrates a computing device that may be used in illustrative
15 implementations of the present invention. With reference to Figure 5, one
16 exemplary system on which the invention may be implemented includes a
17 computing device, such as computing device 500. In a very basic configuration,
18 computing device 500 typically includes at least one processing unit 502 and
19 system memory 504. Depending on the exact configuration and type of
20 computing device, system memory 504 may be volatile (such as RAM), non-
21 volatile (such as ROM, flash memory, etc.) or some combination of the two.
22 System memory 504 typically includes an operating system 505, one or more
23 program modules 506, and may include program data 507. This basic
24 configuration of computing device 500 is illustrated in FIGURE 5 by those
25 components within dashed line 508.

1
2 Computing device 500 may have additional features or functionality. For
3 example, computing device 500 may also include additional data storage devices
4 (removable and/or non-removable) such as, for example, magnetic disks, optical
5 disks, or tape. Such additional storage is illustrated in Figure 5 by removable
6 storage 509 and non-removable storage 510. Computer storage media may
7 include volatile and nonvolatile, removable and non-removable media
8 implemented in any method or technology for storage of information, such as
9 computer readable instructions, data structures, program modules, or other data.
10 System memory 504, removable storage 509 and non-removable storage 510 are
11 all examples of computer storage media. Computer storage media includes, but is
12 not limited to, RAM, ROM, EEPROM, flash memory or other memory
13 technology, CD-ROM, digital versatile disks ("DVD") or other optical storage,
14 magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage
15 devices, or any other medium which can be used to store the desired information
16 and which can be accessed by computing device 500. Any such computer storage
17 media may be part of device 500. Computing device 500 may also have input
18 device(s) 512 such as keyboard 522, mouse 523, pen, voice input device, touch
19 input device, scanner, etc. Output device(s) 514 such as a display, speakers,
20 printer, etc. may also be included. These devices are well known in the art and
21 need not be discussed at length here.

22
23 Computing device 500 may also contain communication connections 516
24 that allow the device to communicate with other computing devices 518, such as
25 over a network. Communication connections 516 is one example of

1 communication media. Communication media may typically be embodied by
2 computer readable instructions, data structures, program modules, or other data in
3 a modulated data signal, such as a carrier wave or other transport mechanism, and
4 includes any information delivery media. The term “modulated data signal”
5 means a signal that has one or more of its characteristics set or changed in such a
6 manner as to encode information in the signal. By way of example, and not
7 limitation, communication media includes wired media such as a wired network or
8 direct-wired connection, and wireless media such as acoustic, RF, infrared and
9 other wireless media. The term computer readable media as used herein includes
10 both storage media and communication media.

11
12 Although details of specific implementations and embodiments are
13 described above, such details are intended to satisfy statutory disclosure
14 obligations rather than to limit the scope of the following claims. Thus, the
15 invention as defined by the claims is not limited to the specific features described
16 above. Rather, the invention is claimed in any of its forms or modifications that
17 fall within the proper scope of the appended claims, appropriately interpreted in
18 accordance with the doctrine of equivalents.